CARBON EMISSIONS AND BUSINESS CYCLES

Hashmat Khan
Christopher R. Knittel
Konstantinos Metaxoglou
Maya Papineau

Carbon Emissions and Business Cycles
Hashmat Khan, Christopher R. Knittel, Konstantinos Metaxoglou, and Maya Papineau
NBER Working Paper No. 22294
May 2016
JEL No. E32,Q43,Q48,Q54

ABSTRACT

U.S. carbon dioxide emissions are highly procyclical—they increase during expansions and fall during recessions. Given this empirical fact, we estimate the response of emissions to four prominent technology shocks from the business-cycle literature using structural vector autoregressive methodologies and data for 1973–2012. By studying the response of emissions to these shocks, we provide a novel approach to assess the shocks' relevance as sources of aggregate output fluctuations. We find that emissions rise on impact only after an anticipated investment-specific technology shock; the response is statistically significant after the first quarter. The same shock explains most— roughly a third—of the total variation in emissions at a horizon of 5 years. Notably, emissions decrease on impact after an unanticipated neutral technology shock in a statistically significant way. This negative empirical response has the opposite sign from its theoretical counterpart in recent environmental DSGE (E-DSGE) models. Since the positive response of emissions drives the E-DSGE models' recommendation for an optimal procyclical policy, our findings suggest that such a policy recommendation should be treated cautiously.

Hashmat Khan
Department of Economics
Carleton University
D891 Loeb, 1125 Colonel By Drive
Ottawa, ON, K1S 5B6
Canada
hashmat.khan@carleton.ca

Christopher R. Knittel
MIT Sloan School of Management
100 Main Street, E62-513
Cambridge, MA 02142
and NBER
knittel@mit.edu

Konstantinos Metaxoglou
Department of Economics
Carleton University
D891 Loeb, 1125 Colonel By Drive
Ottawa, ON, K1S 5B6
Canada
konstantinos.metaxoglou@carleton.ca

Maya Papineau
Department of Economics
Carleton University
D891 Loeb, 1125 Colonel By Drive
Ottawa, ON, K1S 5B6
Canada
maya.papineau@carleton.ca

1 Introduction

U.S. carbon dioxide (CO_2) emissions are highly procyclical. They increase during booms and fall during recessions (Figure 1). The contemporaneous correlation between the cyclical components of GDP and emissions using quarterly data for 1973–2012 is 0.67 and highly significant.[1] This strong comovement between output and emissions in the data suggests that aggregate macroeconomic shocks that drive cyclical fluctuations in output should also account for the procyclicality of emissions.

Identifying the business-cycle shocks that generate procyclical emissions is relevant for both macroeconomics and environmental economics—in particular, the recently developed literature on environmental dynamic stochastic general equilibrium models (henceforth, E-DSGE models). For the macroeconomics literature, it provides a novel way to validate the shocks' relevance as sources of aggregate output fluctuations. To date, such validation exercises are performed mainly using output and hours worked. For the environmental economics literature, shocks that account for the positive output-emissions relationship in the data provide an empirical justification for their use in E-DSGE models that study optimal policy over the business cycle. These E-DSGE models advocate an optimal policy that dampens the procyclicality of emissions based on emissions' positive response to a business-cycle shock. Surprisingly, to our knowledge, there is no applied work that has examined the relationship between business-cycle shocks and emissions. The main contribution of our paper is to fill this void.

Our paper builds on the empirical macroeconomics literature that has identified several structural shocks as potential sources of business cycles. Drawing on this literature, we estimate the response of carbon emissions to four prominent technology shocks—see Ramey (2016)— using quarterly U.S. data between 1973 and 2012. We start with the unanticipated neutral technology (NT) shock of Galí (1999) and Francis and Ramey (2005). This is not only one of the most widely studied business-cycle shocks in the empirical macroeconomics literature, it has also been used in E-DSGE models (see, for example, Fischer and Springborn (2011), Heutel (2012), Annicchiarico and Dio (2015), among others). Additionally, we estimate the response of emissions to an anticipated NT shock (e.g., Barsky and Sims (2011)), an unanticipated investment-specific technology (IST) shock (Fisher (2006)), and an anticipated

---

[1]According to the 2013 Economic Report of the President (ERP), actual 2012 U.S. carbon emissions were approximately 17 percent below their business-as-usual baseline with 52 percent of this reduction attributed to the most recent recession. See the section Carbon Emissions: Progress and Projections on pages 194–195 of the ERP at https://www.whitehouse.gov/administration/eop/cea/economic-report-of-the-President/2013. Doda (2014) documents cross-country evidence for the procyclicality of the emissions.

IST shock (Ben Zeev and Khan (2015)). We identify each shock separately using popular structural vector autoregression (SVAR) specifications and identification restrictions.[2]

We find that emissions *decrease* on impact after a positive unanticipated NT shock in a statistically significant way (Figure 2, panel (a)). The response becomes positive after the third quarter and remains statistically insignificant for horizons up to 5 years. The negative contemporaneous response of emissions is reminiscent of Galí (1999) and Francis and Ramey (2005), who showed a significant decrease in hours worked following a positive unanticipated NT shock triggering a lively debate in the empirical business-cycle literature.[3] Importantly, the empirical contemporaneous response of emissions to the same shock has the *opposite* sign from its theoretical counterpart in recent E-DSGE models. This second observation deserves special attention because the E-DSGE models' recommendation for optimal environmental policy over the business cycle is based on the emissions' positive response to a business-cycle shock.

Emissions decrease on impact after a positive unanticipated IST shock and a positive anticipated NT shock (Figure 2, panels (b) and (c)). These contemporaneous responses are also in sharp contrast with the procyclicality of emissions in the data. In the former case, the emissions' response remains negative for horizons of up to 5 years and it is significant after the first year. In the latter case, the emissions' response becomes positive after two quarters but remains statistically insignificant. The only shock that generates a positive response on impact and for horizons up to 5 years is the anticipated IST shock—the emissions' response is statistically significant after the first quarter (Figure 2, panel (d)).

The anticipated IST shock also accounts for about a third of the emission's variation at a 5-year horizon (Figure 3). The unanticipated NT and IST shocks account for about 18 percent and 20 percent of the variation in emissions at a horizon of 5 years, respectively. The anticipated NT shock accounts for 6 percent of the variation in emissions at the same horizon. Overall, the unanticipated NT shock explains a fairly constant fraction of the variation in emissions across forecasting horizons. While the IST shocks explain a larger fraction of the variation in emissions at longer horizons, the anticipated NT shock does so at shorter horizons.

---

[2]For a recent example on the use of VARs to inform environmental policy, see the report to the California Air Resource Board by Borenstein et al. (2014). Due to the nature of the data, the authors employ a nonlinear VECM methodology to answer two questions related to the the price-collar (floor & ceiling) mechanism on allowance prices for the state's cap-and-trade market for GHG emissions. The first is assessing the probabilities that market prices will be near the floor or the ceiling. The second is the market participants' ability to affect allowance prices using strategic buying and withholding. An executive summary with findings and recommendations is available in pages 2–8 of their report.

[3]Ramey (2016) provides a detailed overview and discussion.

To summarize, we find that, with the exception of the anticipated IST shock, emissions either fall or do not increase on impact in response to four prominent technology shocks (Table 1). In the case of the anticipated IST shock, the positive response becomes statistically significant after the first quarter. This shock also accounts for most—about a third—of the forecast error variance of emissions at a horizon of 5 years. Turning to the other three technology shocks, their failure to generate a positive response of emissions on impact poses a new challenge for establishing their role in driving business cycles. Examining the response of carbon emissions to technology shocks is a novel way to assess their relevance in fluctuations of the economy and complements existing validation exercises using the response of output and hours worked.

Our findings are also relevant for recent work in environmental economics that introduces carbon emissions into calibrated DGSE models to study optimal policy that varies over the business cycle; examples include Annicchiarico and Dio (2015) and Heutel (2012).[4] The main insight offered by these E-DSGE models is that an environmental policy that takes into account the observed procyclicality of emissions is superior to a policy that does not. For example, a key finding in Heutel is that optimal policy—whether a tax or a cap—dampens the procyclicality of emissions. It allows emissions to increase during expansions and decrease during recessions, but not by as much as they would have in the absence of the policy. Since the positive response of emissions to an unanticipated NT shock in the E-DSGE models drives the recommendation for an optimal procyclical policy, our findings suggest that such a policy recommendation should be treated cautiously.

Although anticipated NT and unanticipated IST shocks in an E-DGSE model can generate a procyclical response of emissions—hence, justify procyclical environmental policy—our findings suggest that their empirical counterparts fail to do so.[5] However, incorporating an anticipated IST shock as a business-cycle source in an E-DSGE model will deliver a theoretical response of emissions consistent with its empirical counterpart. Therefore, a research agenda that examines whether environmental policy focusing on the investment sector is able to improve welfare relative to an economy-wide policy is worth pursuing.

---

[4]See also Chang et al. (2009), Angelopoulos et al. (2010), Fischer and Springborn (2011), Dissou and Karnizova (2012), Lintunen and Vilmi (2013), Grodecka and Kuralbayeva (2014). Fischer and Heutel (2013) provides an overview. Annichiarico and Dio consider a consumption shock and a monetary policy shock. Although these non-technology shocks are important for short-run fluctuations, they are not viewed as major sources of business cycles in the macroeconomics literature.

[5]While the existing E-DSGE literature has considered only one type of technology shocks—unanticipated NT shocks—to generate procyclical emissions, we have considered an E-DSGE model to show that anticipated NT and unanticipated IST shocks also generate procyclical emissions. These results are available upon request.

The rest of the paper is organized as follows. In Section 2, we provide an overview of the types of structural shocks we consider along with the identification methodology. In Section 3, we discuss the data and the empirical results. Section 4 concludes. The tables and figures follow the main body of the text.

2   Business Cycle Shocks

2.1   Overview

Given that our objective is to examine the response of carbon emissions to business-cycle shocks, a natural question is which shocks to consider. We draw on the literature that estimates shocks using identification restrictions within an SVAR framework, building on the early work of Shapiro and Watson (1988) and Blanchard and Quah (1989), focusing on technology shocks. In particular, we select four types of technology shocks that are ubiquitous in DSGE models currently used to study business cycles.[6] In the discussion that follows, we provide an overview of each of the four shocks we considered.

i. Unanticipated Neutral Technology (NT) Shocks

In the real business-cycle literature, unanticipated NT shocks that represent exogenous variation in current productivity are the main driver of business cycles. Although the findings of Galí (1999) and Francis and Ramey (2005) suggest that an unanticipated NT shock accounts for only a small variation in output, the recent E-DSGE literature has adopted this type of shock in prescribing environmental policy over the business cycle—see Fischer and Heutel (2013), for a succinct and informative discussion.

ii. Unanticipated Investment-Specific Technology (IST) Shocks

Fluctuations in the price of investment goods relative to the price of consumption goods have also been shown to be important drivers of U.S. business cycles; see Greenwood et al. (2000), Fisher (2006), and, more recently, Justiniano et al. (2010). Exogenous movements in the current relative price of investment goods reflect IST shocks. Although an NT shock affects the production of all goods in a homogenous fashion, an IST shock affects only investment goods.[7]

---

[6]We refer to Ramey (2016) for a detailed discussion of all four technology shocks.

[7]Another type of an investment shock is the marginal efficiency of investment (MEI) shock, which affects the transformation of investment goods into installed capital. In an estimated DSGE model, it is possible to estimate the effects of both investment-specific and MEI shocks; Justiniano et al. find that the MEI shock is

iii. Anticipated NT Shocks

Following Beaudry and Portier (2006), recent work has studied the role of news about a future fundamental in driving business cycles—see Jaimovich and Rebelo (2009), Barsky and Sims (2011)), Khan and Tsoukalas (2012), and Schmitt-Grohe and Uribe (2012).[8] One such fundamental is total factor productivity (TFP) and the notion of anticipated or technology news shocks applies to anticipated movements in future TFP that are uncorrelated with current TFP.

iv. Anticipated IST Shocks

Ben Zeev and Khan (2015) have shown that news about future investment-specific technology is a significant force behind U.S. business cycles. They develop an identification scheme similar to Barsky and Sims (2011) but they focus on the relative price of investment as the fundamental. The measure of investment-specific technical change is the inverse of the relative price of investment representing investment-specific technology. The identification scheme delivers anticipated movements in investment-specific technology that are orthogonal to current investment-specific technology and current TFP.

2.2 Identification

We briefly describe the identification methodology for each of the four shocks discussed in the previous section. Our focus is not on parsing the merits and drawbacks of the empirical approaches but rather take them as standard methods in the business-cycle literature. We consider first-moment technology shocks as they have been extensively considered in the literature.

i. Unanticipated NT Shocks

We follow the methodology in Galí (1999). The identification assumption is that only a technology shock affects labor productivity in the long run. The theoretical rationale for this assumption is that it holds in almost all commonly used business-cycle models. The empirical feasibility of the identification scheme requires a unit root in labor productivity. This is the case for the U.S. and is well documented; see Galí (1999) and Francis and Ramey (2005). Note also that long-run identification using the SVAR approach is equivalent to the instrumental variables (IV) approach in exactly identified systems such as the one considered

---

quantitatively more important than the IST shock using U.S. data. In the SVAR literature, however, these two types of investment shocks have not yet been separately identified.

[8]Beaudry and Portier (2014) provide a detailed overview of this literature.

here.[9] In particular, we consider the following bivariate VAR specification

$$\begin{bmatrix} \Delta LP_t \\ \Delta CO2_t \end{bmatrix} = \begin{bmatrix} C^{11}(L) & C^{12}(L) \\ C^{21}(L) & C^{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^o \end{bmatrix} \equiv C(L)\varepsilon_t = \sum_{j=0}^{\infty} C_j \varepsilon_{t-j}. \tag{1}$$

In terms of notation, ΔLP_t and $\Delta CO2$ are the growth rates of labor productivity and CO_2 emissions per capita, respectively. In addition, ε_t^z is the technology shock to be identified, and ε_t^o is the non-technology shock lacking a structural interpretation. Furthermore, $E[\varepsilon_t \varepsilon_t^{'}] = I$ and $E[\varepsilon_t \varepsilon_s^{'}] = 0$ for $t \neq s$. We use growth rates for carbon emissions because they exhibit a unit root.

We can then state the long-run identification assumption described above as follows

$$C^{12}(1) = \sum_{j=0}^{\infty} C_j^{12} = 0. \tag{2}$$

The reduced-form moving average (MA) representation associated with (1) is

$$\begin{bmatrix} \Delta LP_t \\ \Delta CO2_t \end{bmatrix} = \begin{bmatrix} A^{11}(L) & A^{12}(L) \\ A^{21}(L) & A^{22}(L) \end{bmatrix} \begin{bmatrix} e_{1t} \\ e_{2t} \end{bmatrix} \equiv A(L)e_t = \sum_{j=0}^{\infty} A_j e_{t-j} \tag{3}$$

with $A_0 = I$, $E[e_t e_t^{'}] = \Omega$, $E[e_t e_s^{'}] = 0$ for $t \neq s$, and $\Omega = C_0 C_0^{'}$. In addition, $e_t = C_0 \varepsilon_t$, and $C_j = A_j C_0$. The empirical implementation of (3) proceeds by estimating the following VAR

$$\begin{bmatrix} \Delta LP_t \\ \Delta CO2_t \end{bmatrix} = \begin{bmatrix} B^{11}(L) & B^{12}(L) \\ B^{21}(L) & B^{22}(L) \end{bmatrix} \begin{bmatrix} \Delta LP_{t-1} \\ \Delta CO2_{t-1} \end{bmatrix} + \begin{bmatrix} e_{1t} \\ e_{2t} \end{bmatrix}, \tag{4}$$

where $B^{ij}(L)$ is a polynomial with four lags.

ii. Unanticipated IST Shocks

We follow Fisher (2006) with the key identification assumption being that only an investment

---

[9]See, for example, Shapiro and Watson (1988), Fisher (2006), and, more recently, Section 4.3. in Stock and Watson (2016). In the case of IVs, the econometrician has to confront the possibility of weak instruments in which case a proper methodology for inference is needed, especially when the IVs are highly persistent, or nearly non-stationary—see Chevillon et al. (2015). Chevillon et al. revisit the "hours debate" and show that the first-difference specification in Galí (1999) is not weakly identified.

shock has a long-run effect on the relative price of investment

$$
\begin{bmatrix} \Delta p_t \\ \Delta LP_t \\ \Delta CO2_t \end{bmatrix} = \begin{bmatrix} C^{11}(L) & C^{12}(L) & C^{13}(L) \\ C^{21}(L) & C^{22}(L) & C^{23}(L) \\ C^{31}(L) & C^{32}(L) & C^{33}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^{ist} \\ \varepsilon_t^{z} \\ \varepsilon_t^{o} \end{bmatrix} \equiv C(L)\varepsilon_t = \sum_{j=0}^{\infty} C_j \varepsilon_{t-j}, \qquad (5)
$$

where p_t is the logarithm of the relative price of investment and ε_t^{ist} is the IST shock—the remaining shocks have the same interpretation as in (1). The empirical feasibility of this identification scheme requires a unit root in the relative price of investment. Fisher (2006) shows that this is indeed the case for the U.S. There are two long-run identification assumptions used to estimate the shock. First, only an IST shock affects the relative price of investment. Second, both investment-specific and technology shocks affect labor productivity. These identifying restrictions are implemented in (5) as follows

$$
\begin{aligned}
C^{12}(1) = C^{13}(1) &= 0 & (6) \\
C^{23}(1) &= 0. & (7)
\end{aligned}
$$

iii. Anticipated NT Shocks

We follow the methodology in Barsky and Sims (2011), which is based on the maximum forecast error variance (MFEV) over a medium-run horizon. This medium-run identification approach has advantages relative to long-run restrictions; see, for example, Francis et al. (2014). There are two main identification assumptions. First, it maximizes TFP variation over a medium-run horizon of 10 years. Second, it is orthogonal to innovations in current TFP.

In particular, let y_t be a 5×1 vector of observables with its elements being TFP (TFP_t), real non-durable consumption per capita (C_t), real output per capita (Y_t), CO_2 emissions per capita, and credit spread (CS_t). All variables except for the credit spread are expressed in logarithms. The MA representation of the VAR is

$$
y_t \equiv \begin{bmatrix} TFP_t \\ C_t \\ Y_t \\ CO2_t \\ CS_t \end{bmatrix} = B(L)e_t, \qquad (8)
$$

where $B(L)$ is a matrix polynomial in the lag operator, and e_t is the vector of reduced-

form innovations of appropriate dimension. Furthermore, we assume that there is a linear mapping between the reduced-form innovations e_t and the structural shocks ε_t given by

$$e_t = A\varepsilon_t. \tag{9}$$

Equation (8) and (9) imply a structural MA representation

$$y_t = C(L)\varepsilon_t, \tag{10}$$

where $C(L) \equiv B(L)A$ and $\varepsilon_t = A^{-1}e_t$. The impact matrix A is such that $AA' = \Sigma$, where Σ is the variance-covariance matrix of the reduced-form innovations. Note that there is an infinite number of impact matrices that solve the system. In particular, for some arbitrary orthogonalization, $\widetilde{A}$, the entire space of permissible impact matrices can be written as $\widetilde{A}D$, where D is a an orthonormal matrix.[10] The $h-$step ahead forecast error is

$$y_{t+h} - \widehat{y}_{t+h} = \sum_{\tau=0}^{h} B_\tau \widetilde{A}D\varepsilon_{t+h-\tau}, \tag{11}$$

where B_τ is the matrix of MA coefficients at horizon τ. The contribution of structural shock j to the forecast error variance of variable i is given by

$$\Omega_{i,j} = \sum_{\tau=0}^{h} B_{i,\tau} \widetilde{A}\gamma\gamma'\widetilde{A}'B'_{i,\tau}, \tag{12}$$

where γ is the jth column of D and $B_{i,\tau}$ represents the ith row of the matrix of MA coefficients at horizon τ. We place the current TFP shock in the first position of ε_t and index it as 1. We place the anticipated TFP shock in the second position and index it as 2.

Formally, this identification strategy requires solving the following optimization problem

$$\gamma^* = \text{argmax} \sum_{h=0}^{H} \Omega_{1,2}(h) = \sum_{h=0}^{H}\sum_{\tau=0}^{h} B_{1,\tau} \widetilde{A}\gamma\gamma'\widetilde{A}'B'_{1,\tau} \tag{13}$$

subject to the constraints

$$\widetilde{A}(1,j) = 0 \ \ \forall j > 1 \tag{14}$$
$$\gamma(1,1) = 0 \tag{15}$$
$$\gamma'\gamma = 1. \tag{16}$$

---

[10]We use a Choleski decomposition of the variance-covariance matrix.

The constraints in (14) and (15) ensure that the anticipated NT shock has no contemporaneous effect on TFP. The constraint in (16) is a unit-variance restriction on the identified technology shock.

iv. Anticipated IST Shocks

Following Ben Zeev and Khan (2015), we order TFP and the IST as the first and second variables in a VAR system. We now write (8) as follows

$$
y_t \equiv \begin{bmatrix} TFP_t \\ IST_t \\ Y_t \\ CO2_t \\ CS_t \end{bmatrix} = B(L)e_t,
\tag{17}
$$

where $IST_t = -p_t$, and p_t is the relative price of investment (in logs). We place the current TFP and IST shocks in the first and second positions in the ε_t vector and index them as 1 and 2, respectively. We place the anticipated IST shock in the third position and index it as 3. The identification assumptions are that an anticipated IST shock maximizes the variation in future IST over a medium-term horizon of 10 years and is orthogonal to the innovation in current TFP and current IST. Formally, this identification strategy requires solving the following optimization problem

$$
\gamma^* = \operatorname{argmax} \sum_{h=0}^{H} \Omega_{2,3}(h) \;\; = \;\; \operatorname{argmax} \sum_{h=0}^{H} \sum_{\tau=0}^{h} B_{2,\tau} \widetilde{A} \gamma \gamma' \widetilde{A}' B'_{2,\tau}
\tag{18}
$$

subject to the constraints

$$
\begin{align}
\widetilde{A}(1,j) &= 0 \;\; \forall j > 1 \tag{19} \\
\widetilde{A}(2,j) &= 0 \;\; \forall j > 2 \tag{20} \\
\gamma(1,1) &= 0 \tag{21} \\
\gamma(2,1) &= 0 \tag{22} \\
\gamma'\gamma &= 1. \tag{23}
\end{align}
$$

The first four constraints ensure that the identified anticipated shock has no contemporaneous effect on TFP and IST. Constraint (23) is equivalent to (16).

3 Empirical Analysis

3.1 Data Sources and Preliminary Statistics

The data span the period 1973Q1–2012Q1. We calculated quarterly carbon emissions as follows. First, we obtained monthly total energy CO_2 emissions (million metric tons of carbon dioxide) from Table 12.1 in the Energy Information Administration (EIA) December 2013 Monthly Energy Review. Second, we adjusted these monthly emissions for seasonality using the X-12-ARIMA filter. Finally, we aggregated them to quarterly frequency.[11] We calculate CO_2 emissions per capita using quarterly data on civilian non-institutional population.[12]

The output series is real GDP (billions of chained 2009 dollars) per capita.[13] The labor productivity series is the ratio of real GDP to hours of all persons in the non-farm business sector.[14] The consumption series is real non-durable consumption (billions of chained 2009 dollars) per capita.[15] The quarterly utilization-adjusted TFP series is the difference between the business-sector TFP and the utilization of capital and labor from Fernald (2014), which are available from the Federal Reserve Bank of San Francisco.[16] The relative price of investment is the ratio of the implicit price deflator for nondurable consumption goods to the quality-adjusted investment prices for business equipment & software and consumer durables.[17] Finally, the credit spread is the quarterly average of the difference between the monthly BAA and AAA corporate bond yields, which are available from FRED.

In Figure 1, we plot quarterly emissions and output (GDP) for 1973Q1–2012Q1. In panel (a), the series are measured in per capita terms and are normalized in 1973Q1 levels. Output exhibits a clear upward trend with occasional dips during the recessions increasing by a factor of 1.7 during this period. Emissions exhibit a downward trend until the early 1980s, which is subsequently reversed until the early 2000s, largely in line with the growth of the U.S. economy. The downward trend resumes in the later part of the sample. In general, emissions fall during recessions and tend to fluctuate more than output.

---

[11]See http://www.eia.gov/totalenergy/data/monthly/#environment.

[12]This is the quarterly series CNP16OV from the Federal Reserve Economic Data (FRED).

[13]We construct a quarterly series of real per capita GDP using the FRED series GDPC96, and CNP16OV.

[14]We use the quarterly FRED series GDPC96 and HOANBS, respectively.

[15]We use the quarterly average of the FRED series PCEND and DNDGRG3M086SBEA to obtain real consumption.

[16]See also http://www.frbsf.org/economic-research/total-factor-productivity-tfp/. We use their dtfp_util series, which we transform from percent changes at an annual rate to levels.

[17]We use the quarterly average of the FRED series DNDGRG3M086SBEA. Patrick Higgins from the Federal Reserve Bank of Atlanta generously provided the quarterly series of the quality adjusted investment prices.

In panel (b) of Figure 1, we plot the cyclical components of emissions and output, which we extracted using the H-P filter. The troughs during the recessions can now be identified more easily. The positive comovement of cyclical emissions with cyclical output is clear. The unconditional correlation between emissions and output is 0.67 and highly significant with a p-value less than 0.01 for the entire period— the correlation is even higher (0.83) if we focus on the period before 1985. Consistent with Heutel (2012) and Doda (2014), emissions are cyclically more volatile than GDP. The standard deviation of cyclical emissions is 2.3% compared to 1.6% for output. Prior to 1985, the standard deviation of cyclical emissions is 3% compared to 2% for output.

3.2 Econometric Results

In Figure 2, we plot estimated impulse responses for carbon dioxide emissions along with percentile-based confidence bands following Hall (1992).[18] In all four panels, we consider responses to an orthogonalized shock of size equal to one standard deviation. In Table 1, we list the sign of the impact effect for each of the four types of shocks also indicating whether the effect is statistically significant.

Starting with the unanticipated shocks, panel (a) shows the impulse responses of carbon emissions to a positive unanticipated NT shock. There is a significant decrease on impact after the shock. This negative conditional movement on impact is at odds with the positive unconditional correlation between output and emissions observed in the data. The response turns positive in the third quarter and stabilizes at a positive level; however, it is not statistically significant. In panel (b), we plot the response of emissions to a positive unanticipated IST shock. Emissions do not respond much on impact, as well as for the next couple of quarters, with their responses being negative but small. For horizons exceeding 5 quarters, the responses become statistically significant but remain negative, with a slight decrease in magnitude at longer horizons.

In the case of anticipated shocks, emissions decrease on impact after an anticipated NT shock, as shown in panel (c). The effect becomes positive and exhibits a very modest decline after the first quarter with its long-run values being smaller in magnitude than their contemporaneous counterpart. The responses, however, are not significant, both on impact and for the horizons thereafter. Finally, in panel (d) we plot the response of emissions

---

[18]We follow the terminology in Section D.3. of Lütkepohl (2005) when we refer to the percentile-based confidence bands. We construct "naive" confidence bands obtained by connecting confidence intervals for individual impulse responses.

to a positive anticipated IST shock. Emissions increase on impact after the shock in a hump-shaped manner. Although the response on impact is not statistically significant, the responses for the remaining horizons are.

In Figure 3, we show the share of the emissions' forecast error variance (FEV) explained by each of the four shocks in Figure 2 for horizons up to 5 years. Four clear patterns emerge regarding the share of emissions' variance explained by the four shocks across different forecasting horizons. First, the unanticipated NT shock explains a constant share of the emissions' variance across all horizons. Second, the IST shock explains a larger share at longer horizons. Third, the anticipated NT shock explains a larger share at shorter horizons. Finally, the anticipated IST shock accounts for the largest share among the shocks considered at forecasting horizons of 1 year or longer, explaining roughly a third of the emissions' variance.

In more detail, the unanticipated NT shock accounts for about 18% of the emissions' variance at horizons up to 5 years. The share of the emissions' variance explained by the unanticipated IST shock increases from about 7% at a horizon of 2 years to 20% at a horizon of 5 years. The anticipated IST shock accounts for about 20% of the emissions' variance at a horizon of 1 year and accounts for more than 30% at a horizon of 5 years. Finally, the share of the emissions' variance explained by the anticipated NT shock falls from about 14% at a horizon of 1 quarter to somewhere around 6% at horizons between 2 and 5 years.

To conclude, given the procyclicality of emissions in Figure 1, the impulse responses to all but the anticipated IST shock in Figure 2 are puzzling because a shock that increases output should in theory increase carbon emissions on impact—it is well-documented that output tends to rise after these technology shocks. Overall, the failure of the other three technology shocks to generate a positive response of emissions on impact poses a new challenge for empirical macroeconomists in establishing a particular shock as a major business-cycle driver.

3.3   Robustness Checks

In this section, we check the robustness of our empirical findings using three different exercises. First, we limit our analysis to the period 1985Q1–2012Q1 due to structural changes in the U.S. economy that marked the beginning of the Great Moderation. Second, we use an alternative VAR specification to identify the unanticipated NT shock based on TFP as opposed to labor productivity. Finally, we assume that the source of economic fluctuations are oil-price shocks as opposed to technology shocks.

i. Subsample (1985Q1–2012Q1)

As we already mentioned, the first robustness check for our findings is to limit our sample to the period 1985Q1–2012Q1 largely motivated by the fact that the U.S. economy experienced a series of structural changes in the early 1980s; see, for example, Fernald and Wang (2015) and references therein. Although the shorter sample size makes it harder to get precise estimates for the emissions' responses, the sample is long enough to achieve identification of the shocks.

Based on Figure 4, there are three main points about the response of emissions to the technology shocks for the period 1985Q1–2012Q1 relative to the period 1973Q1–2012Q1 in Figure 2 worth mentioning. First, the negative contemporaneous response to an unanticipated NT shock is still statistically significant. However, the estimated responses at horizons of one year or longer are very close to zero. Second, although the emissions' response to an anticipated NT shock remains statistically insignificant for horizons up to 5 years as in the full sample, the response in the subsample is close to zero at horizons of one year or longer. Finally, the impact response of emissions to an anticipated IST shock is essentially zero, but consistent with the full sample, becomes positive and statistically significant after the first quarter.

According to Figure 5, the general patterns regarding the share of the emissions' forecast-error variance explained by the four types of shocks using the full sample emerge also in the case of the shorter sample. Once again, the unanticipated NT shock explains a constant share of the emissions' variance across all horizons, the IST shocks explain a larger share at longer horizons, while the anticipated NT shock explains a larger share at shorter horizons. In addition, with the exception of the anticipated NT shock, the share of emissions' variance explained by the various shocks considered is generally higher when we restrict our sample to 1985 onward.

In more detail, the share of the emissions' variance explained by the NT shocks is slightly above 25% and exceeds its full-sample counterpart by 8 percentage points. The share of the emissions' variance explained by the unanticipated IST shock increases from 15% at a horizon of 2 years to more than 20% at a horizon of 5 years. The anticipated IST shock explains somewhere between 22% and more than 35% of the emissions' variance at horizons of one year or longer. Finally, the share of the emissions' variance explained by the anticipated NT shock decreases from around 5% at a horizon of one year to 2% at a horizon of 5 years.

ii. Total Factor Productivity

Chang and Hong (2006) suggest using total factor productivity instead of labor productivity with long-run identification of unanticipated NT shocks. Therefore, we checked the robustness of our findings to an alternative specification of (1) replacing the growth in labor productivity with the growth in total factor productivity—see the impulse response function and the FEV decomposition in Figure 6.

Once again, a positive unanticipated NT shock leads to a decrease in emissions. The response remains negative for the next two quarters and then becomes positive beginning with the third quarter at which point it flattens out. This pattern of the emissions' response is similar to the one in panel (a) of Figure 2. The responses fail to be statistically significant both on impact and at subsequent horizons. Moreover, the unanticipated NT shock accounts for about 5% of the emissions' variance for horizons between 1 and 5 years, which is less than a third of the fraction of the emissions' variance explained by the same shock using labor productivity.

iii. Oil Shocks

We also examined the response of emissions to the oil-price shocks in Kilian (2009).[19] Using an SVAR methodology, Kilian identifies oil supply shocks, shocks to global demand for all industrial commodities, and demand shocks that are specific to the global crude oil market interpreted as "precautionary" shocks. Following the author's terminology, we refer to the latter two types of shocks as an aggregate demand shock and oil-market specific demand shock. We estimate the emissions' response to the oil-price shocks modifying the regression analysis in Section III of the paper. In particular, we regress the quarterly growth rate of our per capita carbon dioxide emission series on the three types of oil-price shocks that were readily available for 1975Q1–2007Q4 using up to 20 lags. Following the author's approach, we normalized the shocks such that they raise the price of oil.

In Figure 7, we plot the emissions' response to each of the three types of shocks discussed in the previous paragraph. Emissions respond negatively on impact and at horizons of up to 3 years to a supply disruption (panel (a)). At longer horizons, the responses oscillate around zero, attaining generally small values. Overall, the dynamics associated with a

---

[19]There is a long literature on the macroeconomic effects of oil market shocks with the early work on the topic showing that oil supply shocks have large negative effects on the U.S. economy. However, more recent work has shown that since the 1980s oil shocks have had a smaller impact and much of the variation in oil prices is due to international demand shocks—see Stock and Watson (2016) and the references there in. Stock and Watson find that oil supply shocks explain a very small fraction of the variation in major U.S. macroeconomic variables since the mid 1980s. See their Section 7 on the Macroeconomic Effects of Oil Supply Shocks.

supply disruption fail to generate statistically significant emission responses on impact and at horizons of up to 5 years. An unanticipated demand expansion generates a positive impact response for emissions (panel (b)). The responses remain positive for horizons of up to almost 2 years before they turn negative and bottom out just before the 5-year horizon considered. Finally, emissions respond positively to an oil-market specific unexpected demand shock and decrease rather slowly after the 1st quarter (panel (c)). They reach their trough at a horizon of about 3 years and remain negative for horizons of up to 5 years. The responses to an oil-specific demand shock fail to be significant on impact and throughout the next 5 years.[20]

4 Conclusions

U.S. carbon dioxide emissions are highly procyclical—they increase during expansions and fall during recessions. Given this empirical fact, and drawing on a rich empirical macroeconomics literature, we are the first to estimate the response of emissions to four types of technology shocks that have been identified as sources of business cycles. We consider both anticipated and unanticipated neutral technology (NT) and investment-specific technology (IST) shocks. By studying the response of emissions to these technology shocks, we offer a novel way to assess the shocks' relevance in explaining aggregate fluctuations in the economy that complements existing validation exercises based on the response of output and hours worked.

Using popular SVAR specifications and identification schemes along with quarterly data for 1973–2012, we find that emissions increase on impact only after an anticipated IST stock with the response being significant beyond the first quarter. The same shock also accounts for most–about a third–of the variation in emissions over a 5-year horizon. Importantly, emissions respond negatively to an unanticipated NT shock on impact. This negative empirical response has the opposite sign from its theoretical counterpart in recent environmental

---

[20]We also examined the response of emissions to a government spending shock following Blanchard and Perotti (2002) based on a 5-variable VAR using real government spending, real government tax receipts less transfer payments, output, total real private consumption, and carbon emissions per capita. We employed a Choleski decomposition in which the government spending shock affects the remaining variables in the same quarter but not vice versa. We find that a government spending shock has a negative impact effect on emissions with the emissions' response remaining negative for the next three years. Both the impact response and the responses at longer horizons fail to be statistically significant. The government spending shock explains a very small fraction of the emissions' FEV reaching a plateau of around 3% at horizons of 2–3 quarters. According to Ramey (2011), a concern with the VAR identification scheme in Blanchard and Perotti is that some of what it classifies as shocks to government spending may well be anticipated. She offers a narrative approach (NA) for identifying shocks to government spending arguing that NA shocks capture better the timing of the news about future increases in government spending.

DSGE (E-DSGE) models. Since the positive response of emissions drives the E-DSGE models' recommendation for an optimal cyclical policy, our findings suggest that a cautious treatment of such a policy recommendation is needed. As for the other three technology shocks, their failure to generate a positive response of emissions on impact poses a new challenge for establishing their role in driving business cycles.

Anticipated NT and unanticipated IST shocks driving output fluctuations in E-DSGE models can generate a procyclical response of emissions and, hence, justify procyclical environmental policy. However, our findings suggest that their empirical counterparts will fail to do so. An anticipated IST shock as a business-cycle source in the same class of models, on the other hand, will deliver a theoretical response of emissions consistent with its empirical counterpart. Therefore, examining whether a environmental policy that focuses on the investment sector can improve welfare relative to an economy-wide policy is a research direction for the existing E-DSGE literature to consider.

References

ANGELOPOULOS, K., G. ECONOMIDES, AND A. PHILIPPOPOULOS. 2010. "WHAT IS THE BEST ENVIRONMENTAL POLICY? TAXES, PERMITS, AND RULES UNDER ECONOMICS AND ENVIRONMENTAL UNCERTAINTY." *CESIFO Working Paper No.2980*.

ANNICCHIARICO, B., AND F. D. DIO. 2015. "ENVIRONMENTAL POLICY AND MACROECONOMIC DYNAMICS IN A NEW KEYNESIAN MODEL." *Journal of Environmental Economics and Management*, 69(1): .

BARSKY, R.B., AND E.R. SIMS. 2011. "NEWS SHOCKS AND BUSINESS CYCLES." *Journal of Monetary Economics*, 58 273–289.

BEAUDRY, P., AND F. PORTIER. 2006. "STOCK PRICES, NEWS, AND ECONOMIC FLUCTUATIONS." *American Economic Review*, 96 1293–1307.

BEAUDRY, P., AND F. PORTIER. 2014. "NEWS DRIVEN BUSINESS CYCLES: INSIGHTS AND CHALLENGES." *Journal of Economic Literature*, 52(4) 993–1074.

BEN ZEEV, N., AND H. KHAN. 2015. "INVESTMENT-SPECIFIC NEWS SHOCKS AND U.S. BUSINESS CYCLES." *Journal of Money, Credit and Banking*, 47 1443–1464.

BLANCHARD, O., AND R. PEROTTI. 2002. "AN EMPIRICAL CHARACTERIZATION OF THE DYNAMIC EFFECTS OF CHANGES IN GOVERNMENT SPENDING AND TAXES ON OUTPUT." *Quarterly Journal of Economics*, 117(4): 1329–1368.

BLANCHARD, O., AND D. QUAH. 1989. "THE DYNAMIC EFFECTS OF AGGREGATE DEMAND AND SUPPLY DISTURBANCES." *American Economic Review*, 79(4): 655–673.

BORENSTEIN, S., J. BUSHNELL, F.A. WOLAK, AND M. ZARAGOZA-WATKINS. 2014. "REPORT OF THE MARKET SIMULATION GROUP ON COMPETITIVE SUPPLY/DEMAND BALANCE IN THE CALIFORNIA ALLOWANCE MARKET AND THE POTENTIAL FOR MARKET MANIPULATION." *EI Haas WP 251*.

CHANG, J., J. CHEN, J. SHIEH, AND C. LAI. 2009. "OPTIMAL TAX POLICY, MARKET IMPERFECTIONS, AND ENVIRONMENTAL EXTERNALITIES IN A DYNAMIC OPTIMIZING MACRO MODEL." *Journal of Public Economic Theory*, 11 623–651.

CHANG, Y., AND Y.H. HONG. 2006. "DO TECHNOLOGICAL IMPROVEMENTS IN THE MANUFACTURING SECTOR RAISE OR LOWER EMPLOYMENT?" *American Economic Review*, 96 352–368.

CHEVILLON, G., S. MAVROEIDIS, AND Z. ZHAN. 2015. "ROBUST INFERENCE IN STRUCTURAL VARS WITH LONG-RUN RESTRICTIONS." TECHNICAL REPORT, OXFORD UNIVERSITY.

DISSOU, Y., AND L. KARNIZOVA. 2012. "EMISSION CAP OR EMISSION TAX? A MULTI-SECTOR BUSINESS CYCLE ANALYSIS." *Working Paper, University of Ottawa.*

DODA, B. 2014. "EVIDENCE ON BUSINESS CYCLES AND CO_2 EMISSIONS." *Journal of Macroeconomics*, 40 214–227.

FERNALD, J. G., AND J. CHRISTINA WANG. 2015. "WHY HAS THE CYCLICALITY OF PRODUCTIVITY CHANGED? WHAT DOES IT MEAN?" *Federal Reserve Bank of Boston, Current Policy Perspectives No. 15-6.*

FERNALD, J.G. 2014. "A QUARTERLY-UTILIZATION ADJUSTED SERIES ON TOTAL FACTOR PRODUCTIVITY." *Federal Reserve Bank of San Francisco Working Paper No. 2012-19.*

FISCHER, C., AND G. HEUTEL. 2013. "ENVIRONMENTAL MACROECONOMICS: ENVIRONMENTAL POLICY, BUSINESS CYCLES, AND DIRECTED TECHNICAL CHANGE." *Annual Review of Resource Economics*, 5 197–210.

FISCHER, C., AND M. SPRINGBORN. 2011. "EMISSION TARGETS AND THE REAL BUSINESS CYCLE: INTENSITY TARGETS VERSUS CAPS OR TAXES." *Journal of Environmental Economics and Management*, 62 352–366.

FISHER, J.D.M. 2006. "THE DYNAMIC EFFECTS OF NEUTRAL AND INVESTMENT SPECIFIC TECHNOLOGY SHOCKS." *Journal of Political Economy*, 114 413–451.

FRANCIS, N., M.T. OWYANG, J.E. ROUSH, AND R. DICECIO. 2014. "A FLEXIBLE FINITE-HORIZON ALTERNATIVE TO LONG-RUN RESTRICTIONS WITH AN APPLICATION TO TECHNOLOGY SHOCKS." *Review of Economics and Statistics*, 96 638–647.

FRANCIS, N., AND V. RAMEY. 2005. "IS THE TECHNOLOGY-DRIVEN REAL BUSINESS CYCLE HYPOTHESIS DEAD? SHOCKS AND AGGREGATE FLUCTUATIONS REVISITED." *Journal of Monetary Economics*, 52 1379–1399.

GALÍ, J. 1999. "TECHNOLOGY, EMPLOYMENT AND THE BUSINESS CYCLE: DO TECHNOLOGY SHOCKS EXPLAIN AGGREGATE FLUCTUATIONS?" *American Economic Review*, 89 249–271.

GREENWOOD, J., Z. HERCOWITZ, AND G. HUFFMAN. 2000. "THE ROLE OF INVESTMENT-SPECIFIC TECHNOLOGICAL CHANGE IN THE BUSINESS CYCLE." *European Economic Review*, 44 91–115.

GRODECKA, A., AND K. KURALBAYEVA. 2014. "OPTIMAL ENVIRONMENTAL POLICY, PUBLIC GOODS, AND LABOR MARKETS OVER THE BUSINESS CYCLES." *OxCarre Research Paper 137, University of Oxford.*

HALL, P. 1992. *The Bootstrap and Edgeworth Expansion.*: SPRINGER, NEW YORK.

HEUTEL, G. 2012. "HOW SHOULD ENVIRONMENTAL POLICY RESPOND TO BUSINESS CYCLES? OPTIMAL POLICY UNDER PERSISTENT PRODUCTIVITY SHOCKS." *Review of Economic Dynamics*, 15 244–264.

JAIMOVICH, N., AND S. REBELO. 2009. "CAN NEWS ABOUT THE FUTURE DRIVE THE BUSINESS CYCLE?" *American Economic Review*, 99(4): 1097–1118.

JUSTINIANO, ALEJANDRO, GIORGIO PRIMICERI, AND ANDREA TAMBALOTTI. 2010. "INVESTMENT SHOCKS AND BUSINESS CYCLES." *Journal of Monetary Economics*, 57(1): 132–145.

KHAN, H., AND J. TSOUKALAS. 2012. "THE QUANTITATIVE IMPORTANCE OF NEW SHOCKS IN ESTIMATED DSGE MODELS." *Journal of Money Credit and Banking*, 44 1535–1561.

KILIAN, L. 2009. "NOT ALL OIL PRICE SHOCKS ARE ALIKE: DISENTANGLING DEMAND AND SUPPLY SHOCKS IN THE CRUDE OIL MARKET." *American Economic Review*, 99 1053–1069.

LINTUNEN, J., AND L. VILMI. 2013. "ON OPTIMAL EMISSION CONTROL: TAXES, SUBSTITUTION, AND BUSINESS CYCLES." *Bank of Finland Research Discussion Paper No. 24/2013.*

LÜTKEPOHL, H. 2005. *New Introduction to Multiple Time Series Analysis.*: SPRINGER, NEW YORK.

RAMEY, V. 2016. "MACROECONOMIC SHOCKS AND THEIR PROPAGATION." IN *Handbook of Macroeconomics, Forthcoming.*

RAMEY, V.A. 2011. "IDENTIFYING GOVERNMENT SPENDING SHOCKS: IT'S ALL IN THE TIMING." *Quarterly Journal of Economics*, 126 1–50.

Schmitt-Grohe, S., and M. Uribe. 2012. "What's News in Business Cycles." *Econometrica*, 80 2733–2764.

Shapiro, M.D., and M. Watson. 1988. "Sources of Business Cycle Fluctuations." In *NBER Macroeconomics Annual* . ed. by S. Fischer, 3: MIT Press, 111–156.

Stock, J.H., and M.W. Watson. 2016. "Factor Models and Structural Vector Autoregressions in Macroeconomics." In *Handbook of Macroeconomics, Forthcoming*.

Tables and Figures

Table 1: Contemporaneous empirical impulse responses of emissions to technology shocks

Shock	Impact Effect	Significant?
Unanticipated NT	−	Yes
Unanticipated IST	−	No
Anticipated NT	−	No
Anticipated IST	+	No

Note: We use NT to refer to neutral technology and IST to refers to investment-specific technology. Inference is based on Hall (1992) percentile-based confidence intervals as described in the main text.

Figure 1: Carbon emissions and business cycles
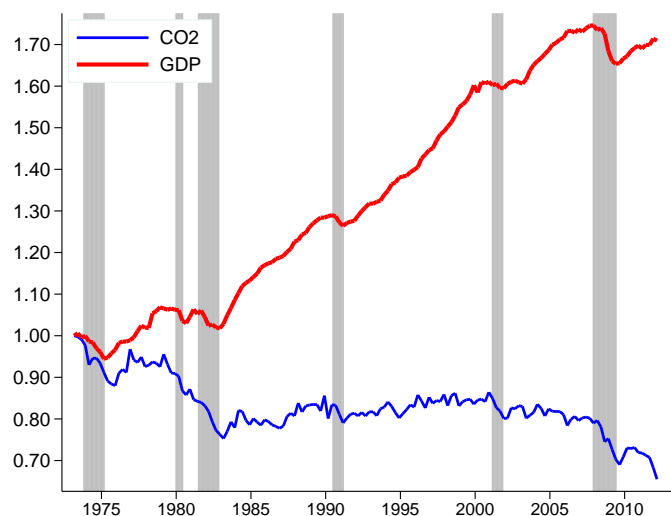
(a) Carbon emissions and GDP normalized to 1973Q1 levels
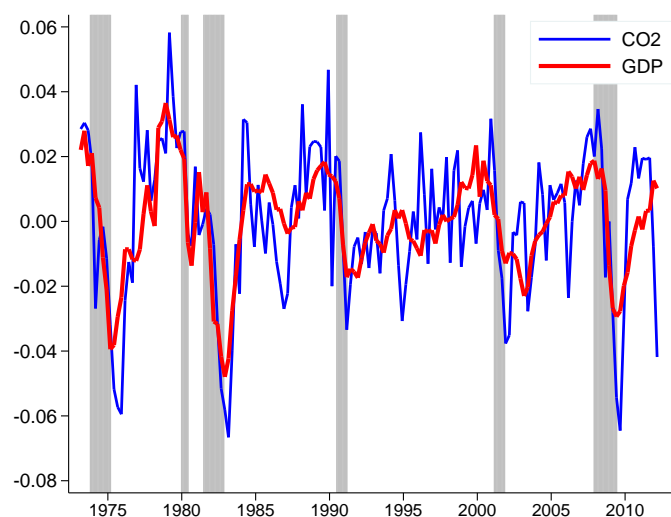
(b) Cyclical components of carbon emissions and GDP

Note: In panel (a), we plot U.S. quarterly real GDP and carbon dioxide emissions for 1973Q1–2012Q1. In panel (b), we plot the cyclical components of log real GDP and log carbon emissions. Both variables are in per capita terms. We extract the cyclical components using the H-P filter. The gray shading identifies NBER recessions.
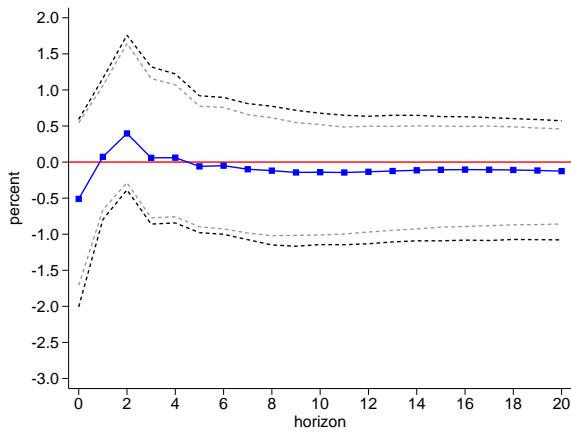
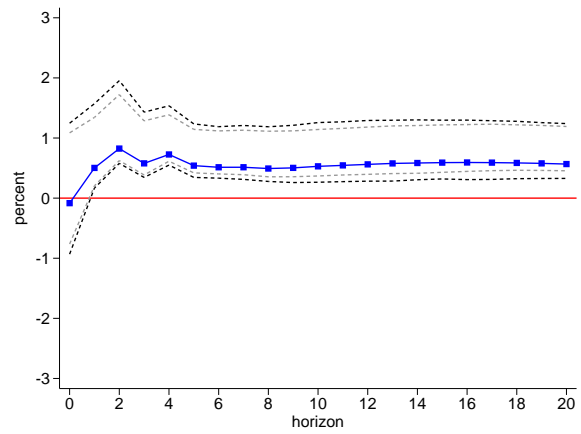Figure 2: Empirical impulse responses of carbon emissions
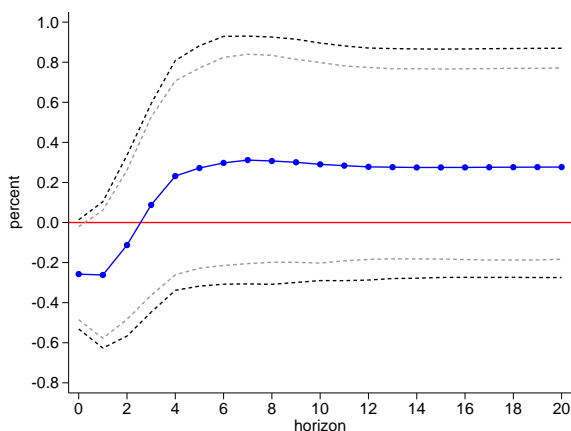
(a) Unanticipated NT shock
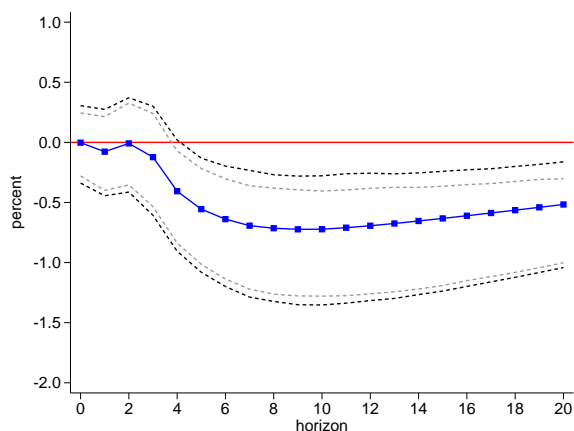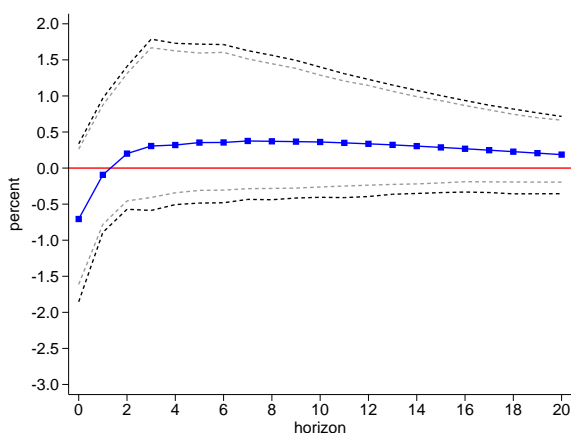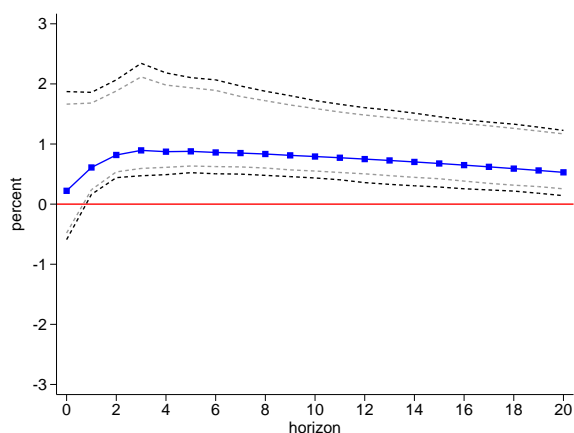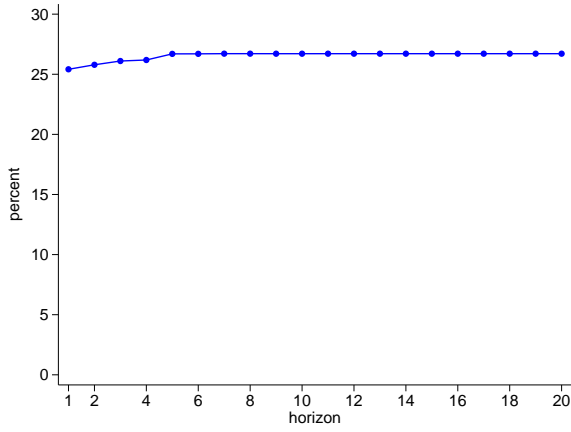
(b) Unanticipated IST shock
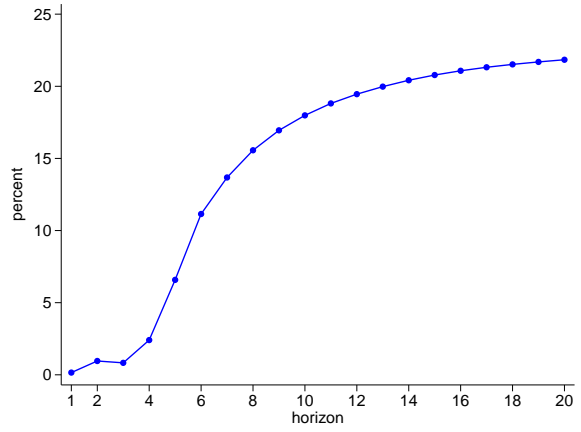
(c) Anticipated NT shock

(d) Anticipated IST shock

Note: We use NT shock to refer to a neutral technology shock. We use IST shock to refer to an investment-specific technology shock. The horizontal axis shows the horizon in quarters. The solid blue lines indicate the empirical impulse responses. The black (grey) short-dashed lines indicate 95% (90%) percentile-based confidence bands constructed using 2,000 bootstrap replications. In all four panels, we consider responses to an orthogonalized shock of size equal to one standard deviation.
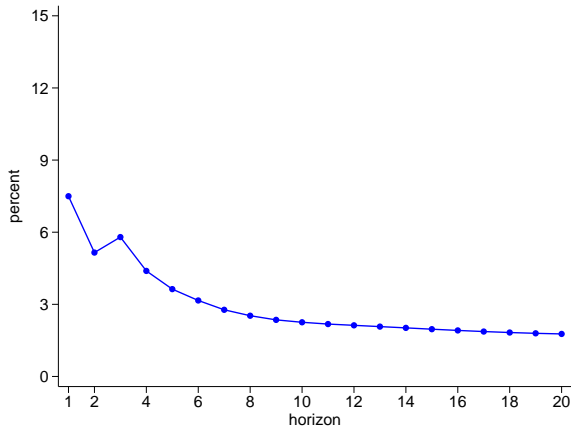
Figure 3: Explained forecast error variance
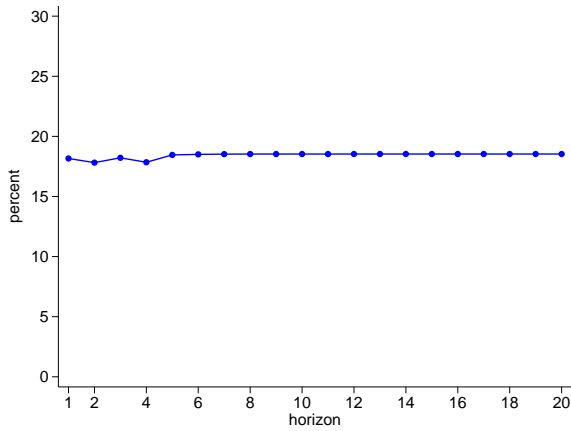
(a) Unanticipated NT shock
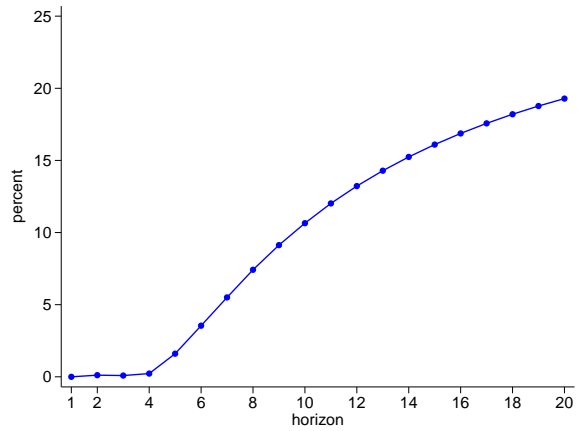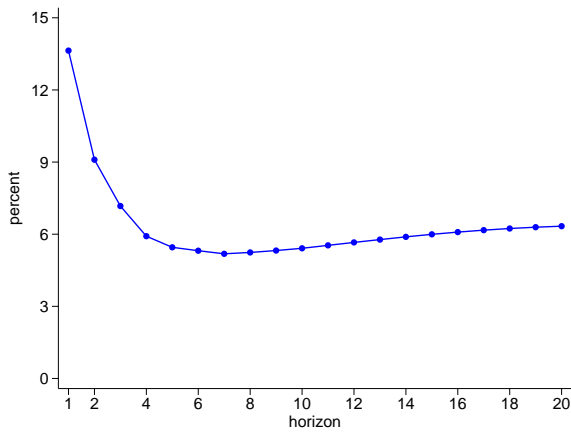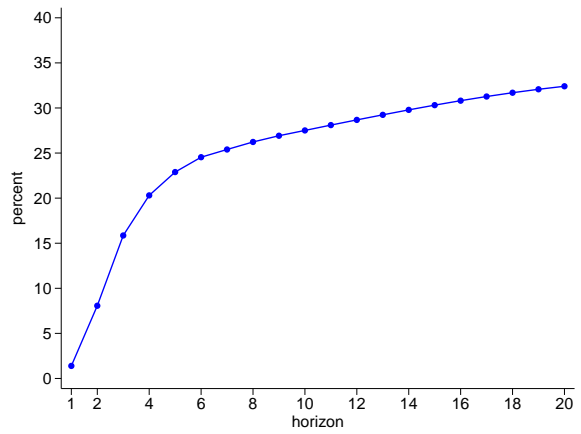
(b) Unanticipated IST shock

(c) Anticipated NT shock

(d) Anticipated IST shock

Note: We use NT shock to refer to a neutral technology shock. We use IST shock to refer to an investment-specific technology shock. The horizontal axis indicates the forecast horizon. The vertical axis measures the fraction of the forecast error variance of carbon emissions explained by each shock.

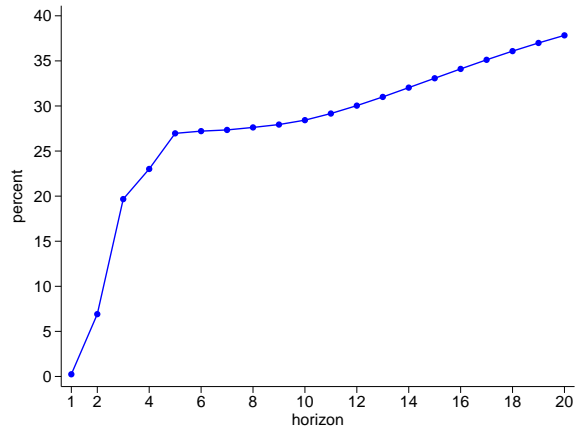Figure 4: Empirical impulse responses of carbon emissions, 1985Q1–2012Q1
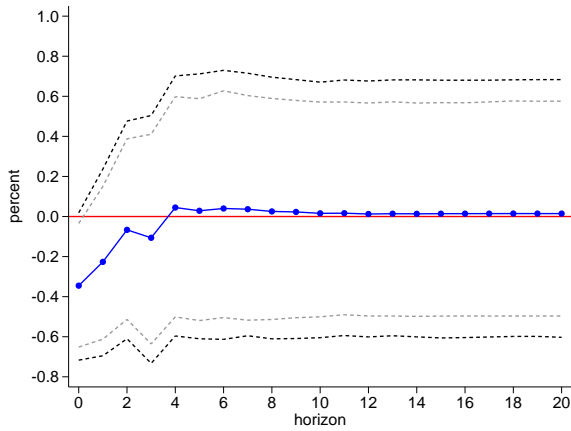
(a) Unanticipated NT shock
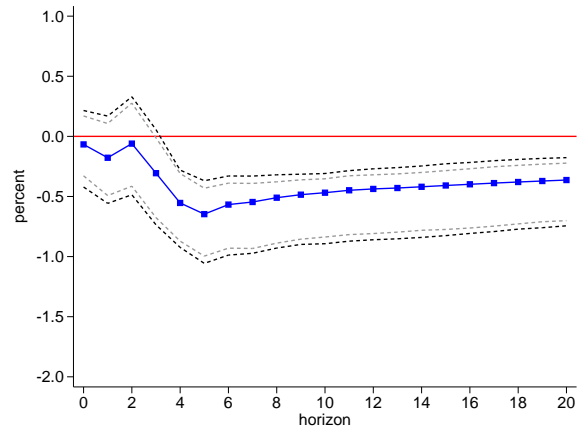
(b) Unanticipated IST shock

(c) Anticipated NT shock

(d) Anticipated IST shock

Note: We use NT shock to refer to a neutral technology shock. We use IST shock to refer to an investment-specific technology shock. The horizontal axis shows the horizon in quarters. The solid blue lines indicate the empirical impulse responses. The black (grey) short-dashed lines indicate 95% (90%) percentile-based confidence bands constructed using 2,000 bootstrap replications. In all four panels, we consider responses to an orthogonalized shock of size equal to one standard deviation.

Figure 5: Explained forecast error variance, 1985Q1–2012Q1

(a) Unanticipated NT shock

(b) Unanticipated IST shock

(c) Anticipated NT shock

(d) Anticipated IST shock

Note: We use NT shock to refer to a neutral technology shock. We use IST shock to refer to an investment-specific technology shock. The horizontal axis indicates the forecast horizon. The vertical axis measures the fraction of the forecast error variance of carbon emissions explained by each shock.

Figure 6: Unanticipated NT shock following Chang and Hong (2006)

(a) Empirical impulse response

(b) Explained forecast error variance
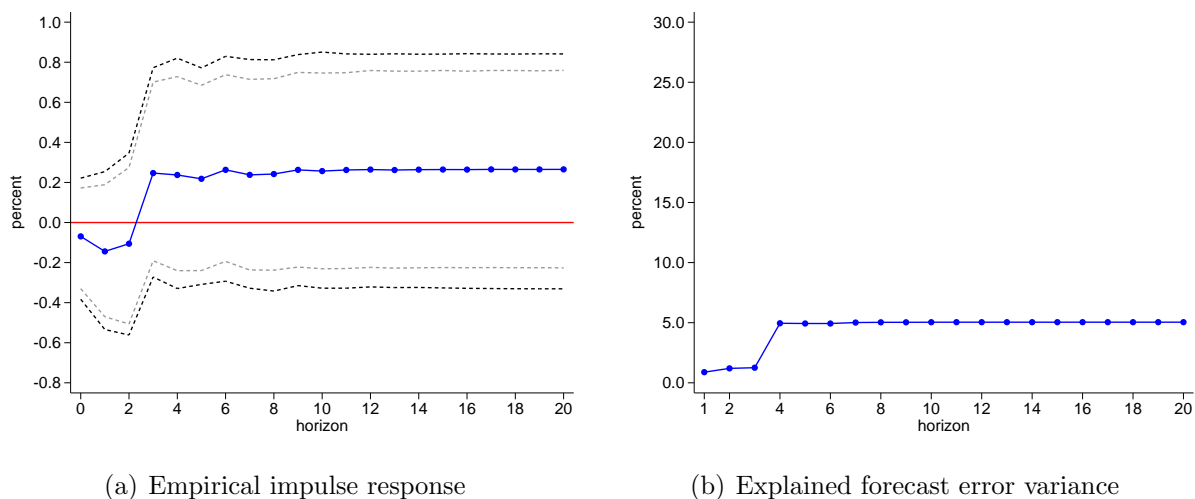
Note: The horizontal axis shows the horizon in quarters. The solid blue lines in panel (a) indicate the empirical impulse responses. The black (grey) short-dashed lines indicate 95% (90%) percentile-based confidence bands constructed using 2,000 bootstrap replications. We consider a response to an orthogonalized shock of size equal to one standard deviation. The vertical axis in panel (b) measures the fraction of the forecast error variance of carbon emissions explained by the unanticipated NT shock.
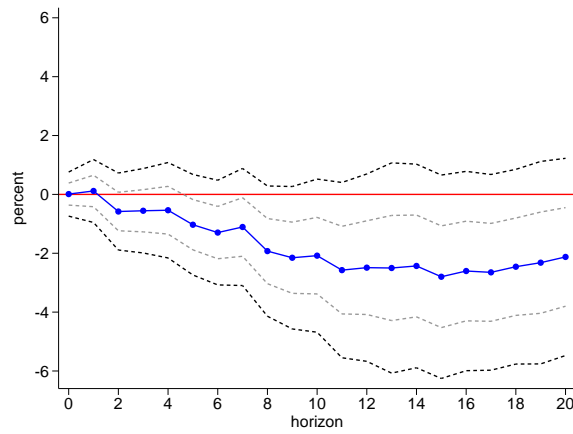
28

Figure 7: Empirical impulse responses of carbon emissions to oil shocks
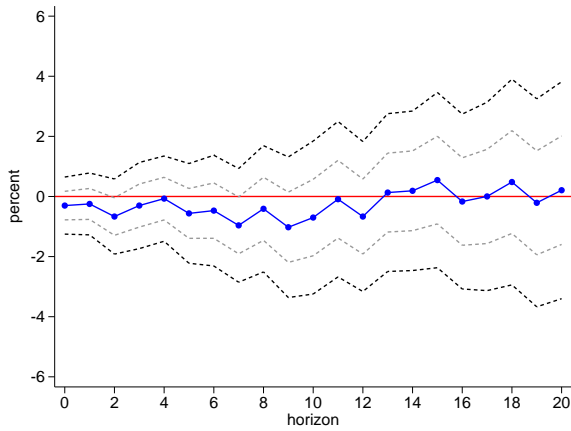
(a) Oil supply shock
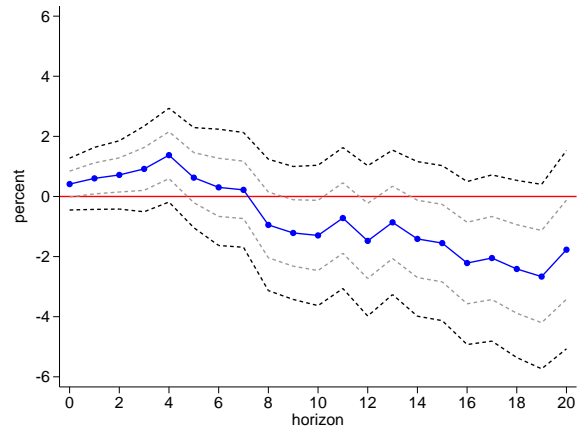
(b) Aggregate demand shock

(c) Oil-market specific demand shock

Note: The horizontal axis shows the horizon in quarters. The solid blue line indicates point estimates. The black (grey) short-dashed lines indicate 2 (1) standard-error confidence bands constructed using the block bootstrap with block size 4 and 20,000 replications as in Kilian (2009).